

Method and Apparatus for Compressing Digital Images

The present invention relates to digital image compression systems. For example, the present invention could be used as part of a CCTV security system, where recorded images need to be compressed to save storage space.

Known image compression systems, such as those utilised in the security industry, suffer accusations of poor image quality in comparison to earlier analogue tape systems, due to the heavy compression techniques used to maximise storage media or use of available bandwidth over networks. In order to increase quality, two methods are known which reduce the amount of stored or transmitted data.

The first method tries to reduce resource utilisation by only saving or transmitting parts of each compressed image that have changed from the last one collected.

Referring to Fig. 1, images from the camera 1 are passed to a digitiser, such as a video decoder 2, where the analogue data is sampled and converted into digital format. This data is then stored in the memory buffer 3. As each image enters the system the incoming image is compared to the last image, or prior images, by

comparator 4. Any differences between the images are identified and this difference data is formatted to match the compression method being used, (for example, multiple 8 by 8 pixel blocks surrounding the required data for JPEG), and then compressed as a subset of the full image, by the compression codec 5. Only this compressed subset is stored and/or transmitted; this can significantly reduce the amount of space required to store the image and/or bandwidth required to transmit it. Images are reconstituted by taking regular "full" images and then adding the subset change data.

However, there are a number of problems associated with this method. Firstly, the final images are temporally distorted, i.e. they are constituted from different images taken at different moments. Secondly, to maintain the integrity of the image, the subsets are compressed at the same level of compression (compression factor) and resolution as the rest of the image, therefore there is a single overall image quality. Finally, this method relies on the accuracy of the change detection to ensure that all potentially valuable image areas are captured. This is often dealt with by ensuring maximum sensitivity to motion, which results in over production of data due to the capture of minor insignificant and "false trigger" motion data.

The second method identifies areas of the image as being of more importance than others and tries to reproduce these at a better quality than the rest of the image.

Referring to Fig. 2, images from the camera 1 are passed to a digitiser, such as a video decoder 2, where the analogue data is sampled and converted into digital format. The image is then compressed ready for storage and/or transmission, by the codec 5. During compression, predetermined areas (regions of interest) that have been identified previously by the user are compressed at a lower level of compression (i.e. less compression) than the rest of the image. The data for each possible input is held by the system software compression driver 6. This generates images that when decompressed have areas that show less compression artefacts, and therefore can display more detail than the areas of lesser interest.

In this method, the complete image is provided for compression at the same image resolution as the identified regions of interest, thus codec time is used for compressing all of the information from areas of the image that have already been identified as being of lesser importance. This method is less efficient than conditional refresh system outlined above because it

generates a relatively large amount of redundant information.

The present invention aims to ameliorate at least some of the above problems. The invention preferably aims to provide a system which maximises the use of limited resources such as storage and/or transmission bandwidth while maintaining good image quality on predetermined areas or areas of the image selected by e.g. a trigger event or any other external stimulus. The method of the invention may therefore be applied to still image compression systems and motion compensated systems.

At its most general, the present invention provides a method of compressing images at resolutions lower than the original image, and at higher levels of compression, but compressing motion areas or regions of interest at higher resolutions and lower levels of compression. The method may be applied to any compression method.

Thus, according to one aspect of the present invention there is provided a method of compressing a digital image stored at a first resolution in a memory, the image including a background region and a region of interest, the method including the steps of:

reading the background region from the memory at a second resolution, the second resolution being lower than the first resolution;

reading the region of interest from the memory at a third resolution, the third resolution being greater than the second resolution;

compressing the background region and the region of interest.

Preferably, compressing the background region and the region of interest includes:

compressing the background region using a first level of compression; and

compressing the region of interest using a second level of compression.

The resolution of an image refers to the detail that can be seen in the image, and is thus a measure of the amount of data in the image. Higher resolution images take up more storage space (e.g. memory space); it is desirable that regions of lesser importance take up less space, hence they may be compressed at lower resolutions. If a region is less important, then it is a waste of compressor time to work on a high resolution form of that region. The present invention saves compressing time by 'throwing away' data for less important regions by reducing the resolution. Thus, less data needs to be compressed for the less important regions. There are a number of ways of reducing the resolution of (i.e. removing data from) a less important region. For

example, the region may be read at a smaller size than the equivalent region in the recorded image, i.e. the region provided for compression is physically smaller. The original image may be e.g. 720 x 288 pixels, whereas e.g. the background region may be provided for compression as e.g. a 360 x 144 pixel image. Such an image may still appear to the eye as a decent representation, but significant detail from the larger image will not be there. In this case, the smaller image contains only about a quarter of the data of the original image, so less time is required to compress it. Alternatively, the background region may be read at less than one quarter of the size of the equivalent region in the recorded image.

Resolution changes differ from the act of compression in two ways. Firstly, data is indiscriminately removed when the resolution is reduced, i.e. valid information may be lost. The overall picture remains recognisable, however; indeed, interpolation may be used to 'smooth' the image. In contrast, compression acts to remove data that is of little use, e.g. data that the eye cannot resolve as well as other data (e.g. high frequency information). Thus, images that are compressed at a low level of compression can appear identical to the original image. The second difference between decimating

an image (i.e. reducing its resolution) and compressing an image is the final format of the image. When an image is decimated, it remains an image that can be seen, whereas when an image is compressed, the processing used means that it becomes a different format (e.g. JPEG) which requires further processing (decompression) before the image can be viewed.

Preferably, the third resolution is the same as the first resolution, i.e. the regions of interest are read at the same detail as they were stored; the region of interest is compressed at its original resolution.

Preferably, reading the background region from the memory includes the steps of:

reading the image from the memory at a second resolution,

masking the region of interest in the read image.

Preferably, the masking step includes blanking out the region of interest with data selected to produce the minimum amount of compressed data for that region. This may be achieved by blacking out the region of interest in the read image. The masking of the regions with black blanks reduces the size of the data from the background region when it is compressed. As mentioned above, the image may be read at a smaller size than the recorded image.

Preferably, the second level of compression is lower than the first level of compression, i.e. the region of interest may be compressed less than the background region (the region of lesser importance).

The region of interest may not be compressed at all; it may be stored at its original resolution.

The image is usually initially recorded as analogue data; the invention may include the step of converting this analogue data into a digital image prior to storing it in the memory.

The invention also provides a method of decompressing an image compressed by the above-described method, the decompressing method including the steps of:

- decompressing the region of interest;

- decompressing the background region;

- reading the decompressed background region and region of interest at a fourth resolution;

- merging the read decompressed region of interest with the read decompressed background region.

Preferably, the method includes displaying a decompressed image.

Preferably, the fourth resolution is greater than the second resolution, so the background region is replicated or interpolated up to e.g. the original dimensions. The decompressed region of interest is then

slotted into place, e.g. into the blanked out areas created by the compressing method.

The fourth resolution may be the same as the third resolution, i.e. the region of interest need not be interpolated.

Preferably, the fourth resolution is the same as the first resolution, so the region of interest maintains its original detail.

The advantages of the above-described methods over known systems include:

the region of interest being at a higher quality in both compression and resolution than the areas of lesser importance, whilst storage and bandwidth utilisation are not greatly increased over the conditional refresh method;

processing and production of redundant information is minimised, because unwanted detail is removed in the reading step;

each image is temporally intact, and therefore even though the quality may be lower in areas of lesser importance, no activity will be missed as the whole image is presented.

This invention thus presents a method that can maximise the use of limited resources, while still giving

the user the quality required to make image data useful, and maintaining temporal integrity.

According to a second aspect of the invention, there is provided apparatus for compressing a digital image stored at a first resolution in a memory, the image including a background region and a region of interest, the apparatus having:

a video compressor/decompressor (codec) for reading the background region and the region of interest from the memory at second and third resolutions respectively, and compressing the background region and the region of interest; and

processing means for controlling the video codec.

Preferably, the background region and the region of interest are compressed at first and second levels of compression respectively.

Preferably, the apparatus include input means for capturing the original image, e.g. the input means may be a CCTV camera.

The input means may capture the image as analogue data; the apparatus may include a digitiser to convert the analogue data to a digital image.

Preferably, the apparatus includes processing means for decompressing the compressed background region and

region of interest, reading them at a fourth resolution, and merging them to form a decompressed image.

Preferably, the apparatus has display means for showing the decompressed image.

The region of interest may be a predetermined area of a particular image (e.g. where any activity is likely to be important), or it may be a region determined by the triggering of a motion detector. For example, a video processor may allow a region of interest to be automatically generated if the requirements of certain parameters are fulfilled.

The apparatus may contain means for setting the parameters used by the video processor in determining whether any regions of interest are present. The means may form a learning system, wherein the means is arranged to note regions in which e.g. movement is usual, and regions in and/or times at which e.g. movement is unusual. If movement is recorded in the unusual period or region, the system may designate it a region of interest.

Embodiments of the invention will now be described in detail with reference to the accompanying drawings, in which:

Fig. 1 is a block diagram representing a known compression apparatus;

Fig. 2 is a block diagram representing a known compression apparatus;

Fig. 3 is a block diagram representing an image compression apparatus which is a first embodiment of the present invention;

Fig. 4 is a block diagram representing an image compression apparatus which is a second embodiment of the invention; and

Fig. 5 is a schematic diagram showing the method of the present invention.

In Fig. 3, images from camera 1 passed to a digitiser, such as video decoder 2, where the analogue data from the camera is sampled and converted into a digital image. The digital image is stored in memory buffer 3. The image is then read by the video codec 5 at a lower resolution than the original image, e.g. less than one quarter of the size of the original. Thus the image has a smaller data size. Usually the image will be no less than one eighth of the size of the original image, otherwise too much information will be lost. Regions of interest in the lower resolution image will be blanked out using e.g. black spaces. This ensures that the minimum amount of space is required to store the compressed version. The codec 5 then compresses the image at a predetermined level of compression e.g. into a

JPEG format. The decompressed background image is formed from this. The region of interest is sent e.g. at the original resolution to the video codec 5, where it is compressed at a predetermined level of compression that may be lower than the level of compression for the background image. The compressing functions and sending of images to the codec 5 is controlled by the collection processor 7. In other words, collection processor 7 determines the resolution at which images are sent to the codec 5 and the level of compressions by which they are compressed. It can be seen therefore that the background region of the image is treated in two different ways to the region of interest. Firstly, the resolution at which it is sent to the codec is reduced - this reduces the amount of data in the region, so some data is removed prior to the compression step, thereby enabling compression to be much more efficient. Secondly, the level of compression at which the background reading is compressed may be higher than that for the region of interest because it is already decided that the background region is less important. Any saving in space thus achieved could be used to improve the image quality of the region of interest, either through the level of compression at which it is compressed, or the resolution at which the codec 5 receives it.

The compressed images may be recorded into a storage medium or decompressed for display.

The display processor 8 is used to control decompression of the stored or transmitted images. When the image is decompressed, the lower resolution image is interpolated up to the same dimensions as the original image, and the decompressed region of interest can be 'slotted' into place. Fig. 5 shows the process in more detail. The image is displayed on display 9 after the decompression and merging of the relevant regions.

Fig. 4 shows a second embodiment of the invention having more than one camera 1. Analogue switch 10 selects the camera to be used, and the video decoder 2 converts the signal from the camera 1 into a digital image to be temporarily stored in the memory buffer 3. The capture controller 11 controls the reading and writing of video images to the memory buffer 3. It also controls the masking of the regions of interest and the separation of data to allow the two-stage compression to proceed. This embodiment has two codecs 5 for allowing the background region and the region of interest to be compressed simultaneously. The designation of regions of interest by the capture controller 11 may be influenced by system parameters 12 set by system soft/firmware 13. The soft/firmware 13 may allow users manually to set the

system parameters, e.g. a fixed area within a particular image, or it may allow the regions of interest to be automatically generated by using a video processor 14, which uses preset parameters to decide whether there are any regions of interest within a particular image. Such a system facilitates the operation of a 'learning' system for defining the regions of interest intelligently, i.e. with a past knowledge of the system.

Fig. 5 shows schematically the steps involved in compressing an image 20 captured at a first resolution e.g. by a camera. When read from e.g. a memory buffer, image 20 is split into a background image 22, which is smaller than the originally recorded image, and a region of interest 21, which is a part of the original image at the same size as it. The part 23 equivalent to the region of interest in the background image is blacked out. Background image 22 looks similar to its larger original, but because it is smaller, it contains less data.

The background image 22 and region of interest 21 are compressed at level of compressions Q_1 and Q_2 respectively into e.g. JPEG format files 24. These files are decompressed into background image 26, which has been interpolated from an image the size of background image 22, and region of interest 25, which (depending on Q_2)

looks similar to the region of interest 21 prior to compression. The background image may not be of high quality due to Q_1 and the interpolation, but this does not matter because it has been designated as less important. In any case, no part of the image is lost, it is just more difficult to resolve objects in the background.

Decompressed region of interest 25 is then slotted into the blanked out portion of decompressed background image 26 to form a final decompressed image 27, which is a mixture of well defined regions of interest together with less clear (but still visible) background objects. All parts of the final decompressed image 27 come from the originally recorded image 20, so the image is temporally intact - no event will be missed, even if it occurs outside the region of interest.

The invention may include any variations, modifications and alternative applications of the above examples, as would be readily apparent to a person skilled in the art, without departing from the scope of the invention in any of its aspects.